


Article

A graph-theoretic approach to understanding emergent behavior in physical systems

Alyssa M. Adams^{1,2*} 

¹ University of Wisconsin-Madison, Department of Bacteriology; amadams4@wisc.edu

² Algorithmic Nature Group, LABORES, Paris, France

* Correspondence: amadams4@wisc.edu

Abstract: The exact dynamics of emergence remains one of the most prominent outstanding questions for the field of complexity science. In this paper, I first discuss different perspectives on emergence in various contexts and then discuss a graph-theoretic representation towards understanding the mechanisms of emergence. In particular, I focus on the importance of defining multiple state spaces for a single system. To test these ideas, I analyze the dynamics of all possible spatial state spaces near the critical temperature in an Ising model. This analysis shows that more deterministic state spaces flip more bits than those that appear more random. More importantly, the rank-order of the deterministic fraction across all spatial state spaces differs around the critical temperature. This indicates a relationship between emergent behavior and state space topology. These results suggest that a graph-theoretic representation of various state spaces may be a fruitful path forward in understanding emergent phenomena in biological systems, which are layered with various state spaces and observational perspectives.

Keywords: Emergence; Ising Model; Information; Computation; State Space

1. Introduction

What are the exact mechanisms of emergence and how can we predict when an emergent phenomenon will happen in a system? For the field of complex systems, this remains one of the most prominent outstanding questions. Given a system of components that interact, is it possible to predict emergent behavior resulting from interacting components and their underlying dynamics without running a complete simulation of the system? The implications of solving this problem are vast for biology, particularly for complex systems and ones that evolve open-endedly.

Emergence is conceptually reminiscent of the idea of innovation, which is a key feature of open-ended systems [1–6]. Open-ended systems are ones that never repeat themselves exactly and continuously innovate in several different ways over time [7]. Innovation broadly encapsulates the appearance of something new, whether it be accessing a new state that was previously unreachable [8,9], or an entirely new state space altogether [10]. The latter case is related to emergence in the sense that a new behavior could occur in a coarse-grained state space.

Over the last few decades, there has been a lot of progress in understanding the concept of open-endedness empirically and theoretically, how it can be measured, and how it could be modeled [2,5,6,11]. Yet, there is still a lack of consensus among researchers if we finally have the means to recreate the strongest form of open-endedness in a model from scratch [12]. This largely rests on a model's ability to continuously reach new states and invent new state spaces. This suggests that the phenomenon of emergence is at the heart of understanding open-ended systems such as the evolution of species, resurgence of pandemics, the development of human technologies, and our impact on climate change, just to name a few.

To address some of these barriers to understanding emergence in a mechanistic sense, I first discusses the relationship between emergence, states and state spaces, computation, and an observer's choice of state space. Since the problem of emergence spans many different fields and sub-fields, one of the goals of this paper is to communicate this topic in



Citation: Adams, A. M. A. graph-theoretic approach to understanding emergent behavior in physical systems. *Preprints* 2021, 1, 0. <https://doi.org/>

Received:

Accepted:

Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

a manner that is not field-specific. Outstanding problems and challenges regarding understanding emergence from several different perspectives is given particular importance in this discussion. The second part of this paper tests some of these ideas in a 2-dimensional Ising model. Results suggest that changes in topological properties of all possible spatial state spaces correspond to the critical temperature at which a phase transition occurs. This phase transition is arguably one form of emergence [13] and this methodology could pave way to understanding mechanisms of emergence for the purposes of predicting emergent behavior.

2. What is emergence?

The exact definition of emergence differs in a variety of ways from one field to another [14]. Colloquially, emergence is understood most famously as the quote from Aristotle: "The whole is different than the sum of the parts." Marletto and Deutch [15] define emergence as: "Specifically, certain sets of collective phenomena can be explained in terms of emergent laws relating them only to each other, without reference to the underlying particles and laws." In other words, emergent phenomena or behaviors are a set of rules that is approximately causally complete within a macrostate description of that system. The set of rules that define the emergence behavior do not rely on lower levels of description. While this definition is useful, the precise definition is still debated among researchers.

Bedau and Chalmers both offer a unique perspective on emergence that categorizes many known examples of emergent phenomena into two distinct categories in [16,17]. The first type of emergence is *weak emergence*, where the behavior of individual entities sums exactly to completely describe the behavior of a group. This type can only be derived from internal dynamics and known external conditions only via simulation [16].

On the other hand, *strong emergence* [16,17] is related to a particular form of innovation in the context of open-ended systems. This kind of innovation is the creation of an entirely new state space that the system performs computations on. In the context of emergence, the emerging behavior cannot be derived from the internal dynamics of the interacting components that form the system. As a result, this is one of the biggest challenges for artificial models to overcome in order to mechanistically model open-ended systems such as living and biologically evolving ones [18]. [19] postulates that the strongest form of open-endedness is one that not only solves problems, but creates new problems for it to solve. In other words, the strongest form of open-endedness would not only be able to compute solutions in a given state space, but invent new state spaces for it to perform computations in.

Given a set of possibilities, our intuition of strong emergence suggests that something totally different and new would have to emerge from a system. In a one-dimensional cellular automata model, for example, our knowledge of the strongest form of open-endedness would suggest this model would need to produce something along the lines of a new cellular automata within it, perhaps with a different neighborhood size or rule types. Or in games like chess and go, a completely new game might emerge from it, and new games would emerge from those, until the things that are emerging that are no longer games but something else entirely [10]. A bottom-up example from biological evolution is the eventual emergence of cells, species, and the eventual Internet that evolved from chemical systems [17,20–22].

But coarse-graining alone is not sufficient to produce new kinds of state spaces that are more dynamically complex than the lower state spaces. Israeli and Goldenfeld explored the relationship between two different state spaces of cellular automata and the complexity of their state trajectories [23]. They explored all possible one-dimensional cellular automata state trajectories (procession of 1D states over time) over all symmetrically-distinct rules. These were spatially coarse-grained according to some function, resulting in spatially smaller cellular automata state trajectories. As a result, the microstate trajectories evolving under more complex rules (according to Wolfram's four complexity classes) could only be

coarse-grained into macrostates that evolve according to *less* complex or similarly complex rules. In other words, coarse-graining could only result in lowering or preserving the complexity class of a state trajectory, not increasing it. Thus coarse-graining a cellular automata pattern made by a less complex rule can never result in a pattern made by a more complex rule.

Biological systems are notorious for their ability to coarse-grain information on almost every level [22]. This would suggest that, at least for bounded, deterministic systems, coarse-graining *decreases* the complexity of the state trajectory. Yet, we understand biological systems as becoming *more* complex over time. This mismatch in expectations could be a result of our ignorance in understanding how multiple state spaces interact in biology, whether it be for multiple state spaces that describe a single system or state spaces of two spatially separate systems that interact through some function. This also affects our ability to understand how computation occurs across many state spaces within a *single* system.

Behavior that is classified as strong emergence cannot be derived from the sum of the dynamics or rules of the individual parts of a system. It is not currently possible to know the future emerging behavior of a group given the current behavior of all its individual components in systems that exhibit strong emergence. To illustrate, the question of strong emergence in viral evolution could be phrased as "Given the current, measurable set of genomic sequences of all known viruses and their hosts, can we predict which new pathogenic viruses might evolve in the next ten years and can we make vaccines now to prevent them?" Given our current understanding of complex systems, evolution, emergence, and open-ended evolution, we cannot.

2.1. A difference in state spaces

A rainbow is an example of an emergent phenomena. This phenomenon not only arises from the interaction of electromagnetic waves refracting through spherical water droplets in the atmosphere, but also the angle at which an observer stands in relation to the cloud vapor *and* the observer's ability to detect some range of electromagnetic frequencies. Thus, this electromagnetic phenomenon depends on interactions occurring in more than one state space (also known as phase space in some fields); one that is unique to spherical water droplets, and another that is unique to the physical characteristics of an observer.

In any example of emergent behavior, the ability to anticipate emergent behavior depends largely on identifying the state space which an observer makes observations and predictions. Without the correct state space, an emergent phenomenon might not be observable. In the example of a rainbow, the position of an observer and the observer's ability to perceive particular electromagnetic wavelengths is crucial to viewing the rainbow. It is worth understanding the relationship between any given state space (assumed to be chosen or defined by an observer) and the ability to detect and make predictions about that emergent phenomena.

One of the reasons that understanding strong emergence is difficult is likely due to addressing multiple different state spaces to describe a single system. In thermodynamics, two of these types of state spaces are often called the "microstate" and the "macrostate." When considering an observer, it may be useful to consider how observers "pick" state spaces that allow making accurate predictions within that state space [24]. Much work has been done in the field of optimal coarse-graining to address the best possible groupings of states into a single state to best describe dynamics [25].

It is possible an observer's choice in a system is reminiscent to compressible descriptions of a system or object, as formalized in algorithmic complexity, also known as Solomonoff-Kolmogorov-Chaitin complexity [26]. From an evolutionary perspective, is it advantageous for observers to select state spaces (descriptions) or a system that are more compressible? Do more compressible state spaces allow an observer to make predictions about their environment more easily? While these questions are not answered directly by this paper, the connections between state spaces, observers, and coarse-graining are adja-

cent to the topic of emergence through the lens of innovation and open-ended biological evolution.

3. Defining a state and a state space

Following these concepts of innovation and emergence together, it may be useful to describe systems from a state-space explicit perspective. Emergence and innovation both evoke concepts related to expressed states, states that are possible and impossible [14, 15], and entirely new state spaces altogether. Here, I suggest a framework that explicitly describes systems in terms of states spaces (sets of all possible states given some variables), the topology of state spaces as a directed graph of states, and how a system expresses and explores these states within that space. Within this framework, it may be easier to understand the mechanisms that drive emergent behavior for the purposes of predicting emergent phenomena before they happen or by only observing the interactions of individual system components without running a simulation of the entire system. However, it remains an open question how states, state spaces, and state space dynamics relate to each other within a single system (or set of systems) that can be described by multiple state spaces at once.

The set of states within a state space (or phase space) is usually defined as all possible configurations of the variables that describe the system. In the Cartesian coordinate system, the set of states is defined by all possible points along its axes. Individual states are identified by their exact position along those coordinates.

In general terms, the set of all possible states within a state space spanned by some set of variables (x, y, z for example) is:

$$\{s_{x,y,z}\} = \{s_{x_1,y_1,z_1}, s_{x_2,y_2,z_2}, \dots, s_{x_L,y_M,z_N}\} \quad (1)$$

for L, M, N possible instances of x, y, z , respectively. If x, y, z are real numbers on a continuous space (like in Cartesian coordinates) then the size of L, M, N are infinite.

A state space is largely dependent on what aspects of a system an observer is able to enumerate [24]. Any system can be described in more than one state space, depending on the observer and what sorts of variables an observer can access. For example, a rock can be described in terms of color, weight, shape, volume, chemical composition, and its ability to look good in a selfie with various Kardashians.

Transitions between states are determined by the rules of the dynamics in that system. The set of transitions between these states defines the topology of the state space, if states are nodes and transitions between them are directed edges. The connectivity of a state space is determined by the dynamical rules that evolve a system over time [27]. If the system is stochastic and any state can transition into any other state in the system in the next time step, then the topology of the state space is fully connected, with directed edges pointing from any state to any other state. Depending on the choice in state space and the underlying dynamics, it is possible to move from any state to any other state in a single time step or given an arbitrary amount of time. In other state spaces, the dynamics may constrain a system to only a few accessible states from its current state [8,9]. Any trajectory of states (succession of states over time) is entirely determined by the topology of the state space, which is determined by the dynamical rules that evolve the system forward in time.

3.1. Computation within a state space

In terms of states, computation is understood as having a mechanistic system (such as a computer) evolving from one state to another state according to some underlying dynamics. The colloquial understanding of computation insists that this process be reliably predictable, such that for any set of input states and underlying dynamics, a predictable set of outcome states is physically realized after the system undergoes transformation under those dynamics.

Under this framework of state spaces, computation can be seen as the ability for a system to move from one state to another state. An observer that performs a computation

in a system must first define a state space to describe that system, and then navigate the possible transitions between states, as defined by the dynamic rules of that system. If a state space is defined as all possible states and the transitions between them, then a computation is a path along a state space's topology from one state to another state.

There has been recent work on interacting state spaces and computation in [28]. Here, a measurement is described as a "copy" from one state space into another. However, because state spaces might use fundamentally different variables to span all possible states within them, a measurement could also be described as a "translate" from one state space into another. In this case, an instrument (or observer) that performs the measurement is used as a function to "translate" the encoding of a state from one state space into a state that exists in a different state space.

4. As experiment using an Ising model

In biology, escape from randomness is seemingly impossible. Even on the smallest scales, randomness causes genetic variance with long-term and macro-sized implications for evolution. On larger scales, random factors such as weather events and asteroids significantly alter the evolutionary trajectory of living systems. As a result, randomness is one of the reasons that making predictions about biology can be so difficult.

To be able to make predictions about biology in a systematic and precise way, we would need to infer deterministic processes in systems that depend on underlying random processes to some degree. This has been done before in quantum physics and thermodynamics. To do the same for biology (and complex systems more broadly), we would need to better understand the relationship between computation as the traversing over states in a state space, and describing a system in multiple different state spaces.

The 2-D Ising model is one of the best-known examples of a phase transition that evolves under pseudo-random rules and is arguably a well-studied case of emergence (in [16], Bedau classifies similar phase transitions as weak emergence). The simple ferromagnetic Ising model described here has a constant external magnetic field with a coupling strength $J = 1$ for nearest-neighbor interactions only. It is closely related to 2-D cellular automata that evolves according to semi-random rules based on a neighborhood size of 4 (adjacent cell neighbors only). The Hamiltonian is given as:

$$H = -J \sum_{\langle i,j \rangle} S_i S_j \quad (2)$$

for spins S_i that can take the values of 1 or -1 .

An example of a configuration of cells (a single state) for a 4×4 version of this model is the following:

1	1	1	-1
1	-1	1	1
1	-1	-1	1
-1	1	1	-1

At each time step, each cell can flip its state (or not) from from a 1 to a -1 , or vice-versa depending on the energy of the whole grid and the difference in energy $dE = -2J \sum_j S_i S_j$ that such a flip would produce. A cell will flip if the resulting energy from that flip would be lower than the current configuration. Flips that result in a higher energy are still allowed, but will only happen with a probability $e^{-dE/T}$, where dE is the change in energy and T is the temperature of the system¹. This ensures a final equilibrium state at the end of each time step.

Around most temperatures, the Ising model will quickly converge to a state of all 1's or all -1 's. But within a specific range of temperatures (near the critical temperature

¹ When viewed as a Markov chain, this is equivalent to Metropolis-Hastings sampling.

$T = 2.269$), the system will not stabilize to a uniform configuration. The emergent behavior, in this case, is the final equilibrium state of the model overall, which results from the individual cells interacting and flipping. Depending on the temperature of the system, this resulting emergent behavior varies and is notably different around the critical temperature. At this temperature, the system goes from being a ferromagnet to a paramagnet. For these reasons, it is an ideal system to explore the relationship between its emergent properties and the topological properties of the various state spaces to encode the system. A full description of an Ising model and its properties can be found at [29].

5. Materials and Methods

All possible combinations of cells are considered as spatial state spaces, not accounting for rotational symmetries. For a 3×3 grid Ising model, there are 512 distinct possible sets of one or more cells, corresponding to 512 spatial state spaces. Larger Ising Models can be studied by randomly sampling state spaces within the model, since the number of possible state spaces grows exponentially with model size. Only the 3×3 case is studied in this paper as an exploratory case.

Monte Carlo simulations for the 3×3 Ising model were ran for 1000 time steps at several temperatures between the values of 2 and 3, since this is around the critical temperature of a 2-D Ising model. The model was simulated 100 different times for each temperature. For each of these, the following values were calculated for each of the possible spatial state spaces:

- Magnetization at each step (M)
- Average magnetization over the state trajectory of 1000 steps ($\langle M \rangle$)
- Number of bits (cells) flipped at each step normalized by the size of the state space
- Average number of bits flipped over the state trajectory of 1000 steps
- Total number of bits flipped over the state trajectory of 1000 steps

At every step, the magnetization of a state is defined as the absolute value of the sum of all 1's and -1's in the state, normalized by the size of the state:

$$M = \frac{|\sum 1's + \sum -1's|}{N} \quad (3)$$

where N is the number of cells in that state.

For each state space (some subset of cells in the system), the succession of expressed states was tracked over time. From this, the apparent topology of each state space was approximated by representing the transitions between states in a graph. In these graphs, nodes are states within that state space and directed edges denote transitions between the states. As an example, consider the state space of the single upper-left cell of the 3×3 grid. There are only two possible states in this state space, 1 and -1. For the state space that consists of the top row of the grid, there are $2^3 = 8$ possible states. For any state space, there are A^N possible states, where A is the number of possible spins and N is the size of the state space.

A state space that behaves completely randomly is equivalent to a Markov Chain with the transition probabilities between states being equally likely, and all transitions between states are allowed. In a state space that is completely deterministic, the transitions between different states are not probabilistic and the next state can be determined by the current state of the system. In this case, the current state would only have one possible state it could move to, whether it be a new state (making it a transient state) or sitting on the current state indefinitely as an absorbing state². In the graph representation of a completely deterministic state space, each state (node) would only have a single out-degree (the number of edges that points away from the current node to another node). In a state

² However, for this analysis, self-loops (edges that point from the current node back to itself) are not counted.

space that is entirely stochastic, every node would have an out-degree that corresponds to the number of total states in the state space, minus 1 if not counting self-loops.

The "amount of determinism" (here, called the **deterministic fraction**) in a state space graph can arguably be estimated by counting the fraction of nodes with more than one out-degree to the total number of nodes (states) in the state space:

$$\text{deterministic fraction} = \frac{\# \text{ nodes with out degree} > 1}{\# \text{ nodes}} \quad (4)$$

A state space graph that is completely random will have all nodes with an out-degree higher than 1. In contrast, a completely deterministic topology will only have nodes with an out-degree of exactly 1 (or 0 if it is an absorbing state). Since each state space graph is strictly built from the trajectory of states, states that are not expressed within a single trajectory are not counted for this measure.

Because several Python repositories already implement a Monte Carlo simulation of a 2-D Ising model, I started with this Ising model found at rajeshrinet.github.io [30]. Extensive modifications to the code were made to accommodate all possible spatial state spaces and to make the above measurements within them. All Ising model simulations were performed in Python 3.7 and all code is publicly available on GitHub via [alyssa-adams/ising_statespaces](https://github.com/alyssa-adams/ising_statespaces). Intermediate Python pickle files used to store data and results of calculations are also publicly available on [Google Drive](https://drive.google.com).

6. Results

For each figure in this section, the colors consistently denote different temperatures, with higher temperatures in a darker shade. Box plots are shown as averages over 100 runs and scatter plots are shown as the aggregation over all 100 runs. All boxes denote a 95% confidence interval with outliers marked in black.

6.1. Average Magnetization

To validate the Ising model simulation, the average magnetization over 1000 steps was calculated at various temperatures for 100 simulations. These results are shown in Figure 1, which shows the expected behavior for this Ising model in the last column (state space size = 9). The same calculation was made over all other possible state space sizes in Figure 1, shown as columns. As a result, the average magnetization decreases slightly as a function of state space size. State spaces consisting of a single cell (size 1, shown on the left) only have one bit and are always fully "magnetized", but since these state spaces are single cells, this is an artifact of the measure and carries little interpretable meaning.




Figure 1. Average magnetization of an Ising model over all possible spatial state spaces. Results are grouped according to the size of the state space.

6.2. Deterministic Fraction

In addition to looking at average values over 1000 time steps, I measured the deterministic fraction (equation 4) of each state space topology, as defined in Section 5. Figure 2 shows these values grouped according to state space size for various temperatures. The deterministic fraction is always 1 for a state space the size of a single cell, since self-loops are not counted. It is always 0 for state spaces two cells large, but increases with the size of the state space. The differences between temperatures decreases with size as well as variance within individual temperatures.




Figure 2. Deterministic fraction of all state space topologies when grouped by the size of the state space.

Regardless of state space size, the rank-order distribution of deterministic fraction across all state spaces remains constant over all 100 simulations for a given temperature (Figure 3). In this figure, error bars are shown but are small. Increasing the temperature correlates to the deterministic fraction distribution decreasing overall. There is a noticeable difference in the distribution around the critical temperature, above it, and below it.




Figure 3. Rank-order distribution of the deterministic fraction of all state spaces for various temperatures. The size of the individual state spaces is not denoted and the ordering of state spaces is not preserved across the 100 simulations. Error bars showing a 95% confidence interval are shaded. The critical temperature $T_c = 2.269$ is denoted with a dotted line.

The values from Figure 3 are plotted against the total number of bits flipped for each state space trajectory, which is *not normalized* to the size of the state space in Figure 4. For lower temperatures, the relationship appears to be linear, but less so for higher temperatures. In general, **more deterministic state space topologies result in more total bits flipped over a state trajectory.**




Figure 4. The deterministic fraction of all state spaces for versus the total number of bits flipped for all state spaces. Data points are shown for all 100 simulations as an aggregate.

Similar to Figure 3, the ranked order distribution of the total number of bits flipped in a state trajectory (not normalized to state space size) is shown in Figure 5. Step-like distributions are an artifact of having discrete-sized state spaces.

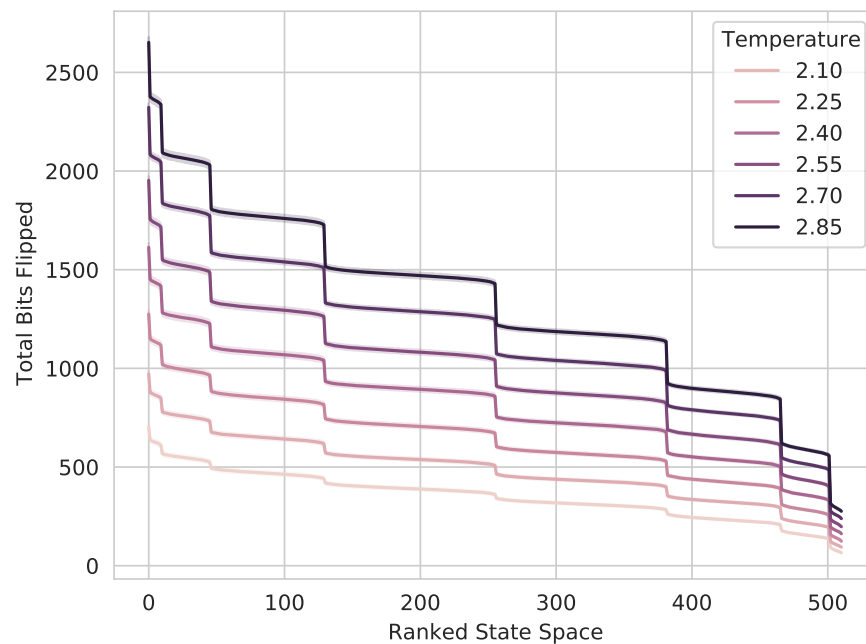


Figure 5. Rank-order distribution of the total number of bits flipped for state trajectories of all state spaces for various temperatures. The size of the individual state spaces is not denoted and the ordering of state spaces is not preserved across the 100 simulations. Error bars showing a 95% confidence interval are shaded.

7. Discussion

Based on these results, the topologies of all possible spatial state spaces of a ferromagnetic 2-D Ising model are correlated with the emergent behavior of the Ising model overall. This result is captured in 3. The emergent paramagnetic behavior from interactions between individual cells around the critical temperature is correlated with a change in the rank-order distribution of the deterministic fraction across all possible spatial state spaces. However, it remains an open question whether emergent behavior can be determined without running the entire simulation. Since the state space graphs are constructed from simulating the model, it remains an open question whether this approach can predict emergent behavior in only a small number of time steps, rather than the full 1000 time steps. In other words, how long can the simulation be run before the the results in 3 are apparent? This is also known as the Markov chain mixing time for this system.

Figure 2 suggests that the cumulative topologies (including all nodes with multiple out-degrees) of state spaces vary in their deterministic qualities. Smaller state spaces are generally less deterministic and larger state spaces are more deterministic overall. This could be due to the fact that 1000 time steps may not be long enough for a state space to explore all possible states. For the largest state space of 9 cells, there are $2^9 = 512$ possible states to choose from and the system may not have had sufficient enough time to explore the possible state transitions that smaller state spaces do. However, in biology, most larger systems do not have sufficient time to explore all possible states and are often forced to evolve regardless of the fraction of states left unexplored. These results are reminiscent of this same phenomenon in biological systems. Smaller systems have less states and are able to explore all possible states and their transitions more thoroughly than larger ones. This could be a driving factor in the evolution of open-ended systems and their ability to compute and solve problems.

Regardless of the size of the state space, the distribution of the deterministic fraction over all state space topologies in Figure 3 remains constant over multiple simulations. The

distribution only changes as a function of temperature; lower temperatures result in more deterministic topologies.

Figure 4 shows the relationship between the total number of bits flipped, the deterministic fraction, and the temperature. Based on these results, it is surprising that the number of bits flips generally increases with the deterministic fraction of the state space topology. In other words, **topologies that are less random and more deterministic flip more bits and carry the energetic burden of computation**. This is surprising because an increase in randomness is often associated with more change. However, there are two different notions of randomness here. The increase in randomness of bits flipping introduced by higher temperatures does not necessarily translate to randomness in the state space topology. Instead, Figure 4 suggests the opposite.

Larger Ising models can be studied in the same way, though would be computationally taxing to do exhaustively as for the 3×3 case. State spaces can be randomly sampled for larger Ising models. It would also be worth exploring these results in future work by breaking down the trajectory of 1000 steps into the parts before the phase transition and after the phase transition. This would illuminate more on the dynamics of the state space before, during, and after an emergent event occurs.

Overall, these results suggest that a graph-theoretic approach to state spaces within a single system may be a productive way to mechanistically capture and predict strong emergence withing a dynamical system.

Funding: This research received no external funding.

Data Availability Statement: All Ising model simulations were performed in Python 3.7 and all code is publicly available on GitHub via [alyssa-adams/ising_statespaces](https://github.com/alyssa-adams/ising_statespaces). Intermediate pickle files used to store data and results of calculations are also publicly available on [Google Drive](#).

Acknowledgments: Special thanks to Joseph L. Austerweil for providing computational resources to run this experiment, lively discussions, suggestions for the flow of ideas, and helpful critiques and edits.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Bedau, M.A.; McCaskill, J.S.; Packard, N.H.; Rasmussen, S.; Adami, C.; Green, D.G.; Ikegami, T.; Kaneko, K.; Ray, T.S. Open problems in artificial life. *Artificial life* **2000**, *6*, 363–376.
2. Banzhaf, W.; Baumgaertner, B.; Beslon, G.; Doursat, R.; Foster, J.A.; McMullin, B.; De Melo, V.V.; Miconi, T.; Spector, L.; Stepney, S.; others. Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences* **2016**, pp. 1–31.
3. Ruiz-Mirazo, K.; Umerez, J.; Moreno, A. Enabling conditions for open-ended evolution. *Biology & Philosophy* **2008**, *23*, 67–85.
4. Taylor, T.J. From artificial evolution to artificial life. PhD thesis, University of Edinburgh. College of Science and Engineering. School of Informatics., 1999.
5. Taylor, T.; Bedau, M.; Channon, A.; Ackley, D.; Banzhaf, W.; Beslon, G.; Dolson, E.; Froese, T.; Hickinbotham, S.; Ikegami, T.; others. Open-Ended Evolution: Perspectives from the OEE1 Workshop in York. *Artificial Life* **2016**.
6. Packard, N.; Bedau, M.A.; Channon, A.; Ikegami, T.; Rasmussen, S.; Stanley, K.O.; Taylor, T. An Overview of Open-Ended Evolution: Editorial Introduction to the Open-Ended Evolution II Special Issue. *Artificial Life* **2019**, *25*, 93–103, [https://direct.mit.edu/artl/article-pdf/25/2/93/1896734/artl_a_00291.pdf]. doi:10.1162/artl_a_00291.
7. Dolson, E.; Vostinar, A.; Wiser, M.; Ofria, C. The MODES Toolbox: Measurements of Open-Ended Dynamics in Evolving Systems. *Artificial Life* **2019**, *25*, 50–73. doi:10.1162/artl_a_00280.
8. Kauffman, S.; Logan, R.; Este, R.; Hobill, D.; Shmulevich, I. Propagating organization: An enquiry. *Biology and Philosophy* **2008**, *23*, 27–45. doi:10.1007/s10539-007-9066-x.
9. Kauffman, S. At Home in the Universe: The Search for the Laws of Self-Organization and Complexity, 1995.
10. Davies, P. *The Demon in the Machine: How Hidden Webs of Information Are Finally Solving the Mystery of Life*; Penguin Books Limited, 2019.
11. Pattee, H.H.; Sayama, H. Evolved Open-Endedness, Not Open-Ended Evolution. *Artificial Life* **2019**, *25*, 4–8, [https://direct.mit.edu/artl/article-pdf/25/1/4/1667089/artl_a_00276.pdf]. doi:10.1162/artl_a_00276.
12. Hintze, A. Open-Endedness for the Sake of Open-Endedness. *Artificial Life* **2019**, *25*, 198–206, [https://direct.mit.edu/artl/article-pdf/25/2/198/1896695/artl_a_00289.pdf]. doi:10.1162/artl_a_00289.
13. Entropy, 2020.

14. Clayton, P.; Davies, P. *The Re-Emergence of Emergence: The Emergentist Hypothesis From Science to Religion*; Oxford University Press, 2006.
15. Deutsch, D. Constructor theory. *Synthese* **2013**, *190*, 4331–4359.
16. Bedau, M.A. Weak Emergence. *Noûs* **1997**, *31*, 375–399, [<https://onlinelibrary.wiley.com/doi/pdf/10.1111/0029-4624.31.s11.17>]. doi:<https://doi.org/10.1111/0029-4624.31.s11.17>.
17. Chalmers, D.J. Strong and Weak Emergence. In *The Re-Emergence of Emergence: The Emergentist Hypothesis From Science to Religion*; Davies, P.; Clayton, P., Eds.; Oxford University Press, 2006.
18. Adams, A.; Zenil, H.; Davies, P.; Walker, S. Formal Definitions of Unbounded Evolution and Innovation Reveal Universal Mechanisms for Open-Ended Evolution in Dynamical Systems. *Scientific Reports* **2017**, *7*, 997.
19. Stanley, K.; Lehman, J.; Soros, L. Open-endedness: The last grand challenge you've never heard of, 2017.
20. Walker, S.I. Top-down causation and the rise of information in the emergence of life. *Information* **2014**, *5*, 424–439.
21. Walker, S.I.; Davies, P.C. The algorithmic origins of life. *Journal of the Royal Society Interface* **2013**, *10*, 20120869.
22. Walker, S.I.; Cisneros, L.; Davies, P.C. Evolutionary transitions and top-down causation. *Proceedings of Artificial Life XIII* **2012**, pp. 283–290.
23. Israeli, N.; Goldenfeld, N. Coarse-graining of cellular automata, emergence, and the predictability of complex systems. *Phys. Rev. E* **2006**, *73*, 026203. doi:10.1103/PhysRevE.73.026203.
24. Longo, G.; Montévil, M. Extended criticality, phase spaces and enablement in biology. *Chaos, Solitons & Fractals* **2013**, *55*, 64–79. Emergent Critical Brain Dynamics, doi:<https://doi.org/10.1016/j.chaos.2013.03.008>.
25. Shalizi, C.; Moore, C. What Is a Macrostate? Subjective Observations and Objective Dynamics, 2003.
26. Soler-Toscano, F.; Zenil, H.; Delahaye, J.P.; Gauvrit, N. Calculating Kolmogorov Complexity from the Output Frequency Distributions of Small Turing Machines. *PLOS ONE* **2014**, *9*, 1–18. doi:10.1371/journal.pone.0096223.
27. Adams, A.M.; Berner, A.; Davies, P.C.W.; Walker, S.I. Physical Universality, State-Dependent Dynamical Laws and Open-Ended Novelty. *Entropy* **2017**, *19*. doi:10.3390/e19090461.
28. Kolchinsky, A.; Corominas-Murtra, B. Decomposing information into copying versus transformation. *J. R. Soc. Interface* **2020**, *17*.
29. Bhattacharjee, S.M.; Khare, A. Fifty years of the exact solution of the two-dimensional Ising model by Onsager. *Current Science* **1995**, *69*, 816–821.
30. Singh, R. Ising model. *GitHub.io* **2014**.